

# Grafana at MWT2 and OSG Connect

Lincoln Bryant  
[lincolnb@hep.uchicago.edu](mailto:lincolnb@hep.uchicago.edu)  
June 9, 2015

# Some background

- Used Cacti, Ganglia, and other tools in the past.
- Ganglia felt too inflexible, Cacti too convoluted
- Heard about Graphite at a Linux User Group presentation, started experimenting in summer 2014
- Started using Grafana around January 2015, very pleased with it thus far.

# Graphite

- <http://graphite.wikidot.com/>
- Real-time graphing for time series data
- Three main components:
  - Whisper - Data format, replacement for RRD
  - Carbon - Listener daemon
  - Front-end - Web interface

# The Graphite Protocol

- Dead simple.
- Open a socket, fire off metrics in the form of:
  - `path.to.metric <value> <timestamp>`

# Sample script

```
#!/bin/bash
metric="htcondor.running"
value=$(condor_q | grep R | wc -l)
timestamp=$(date +%s)
echo "$metric $value $timestamp" | nc \
  graphite.yourdomain.edu 2003
```

# Enter Grafana

- Graphite web interface is powerful but somewhat lacking
- Fortunately it can export JSON instead of PNGs
- Grafana is another tool that sits atop of Graphite
  - provides much nicer plotting tools, dashboarding, etc
- All HTML5 / Javascript goodness.
  - Excellent options for embedding graphs elsewhere.

# Another handy tool - CollectD

- Graphite is almost too simple sometimes.
- We're starting to heavily invest in CollectD
  - Generic system statistics collector
  - Many, many plugins already available, including Graphite
  - Straight-forward to write your own collectors.


# A Gallery of Grafana Graphs




# ATLAS batch queue data

- Scrapes ATLAS Production & Distributed Analysis (PanDA) system via JSON API
- e.g., [http://bigpanda.cern.ch/jobs/?jobtype=analysis&display\\_limit=100&computingelement=MW\\_T2\\_SL6](http://bigpanda.cern.ch/jobs/?jobtype=analysis&display_limit=100&computingelement=MW_T2_SL6)
- Once an hour, as to not overload their webserver
- Breakdowns by sub-site (IU, UC, UIUC)

# ATLAS analysis queue




# ATLAS production queue




# dCache SRM

- Monitors SRM/GridFTP transfers at MWT2
- Scraping dCache's PostgreSQL database
- Runs every 5 min.
- SRMWatch now deprecated -- very suitable replacement.

# dCache SRM







# dCache SRM



# UChicago Tier 3 XRootD

- Scraping network statistics of XRootD servers via CollectD
- Also have an XRootD cache server for caching data from the ATLAS data federation (FAX).
- Looking at cache occupancy, cached file ages, percentage of file cached

# UChicago Tier 3 XRootD



Cache disk occupancy [last 24h](#)

**1.772 TiB**

Age of oldest file in cache [last 24h](#)

**45 days**

16

Average percentage of file cached [last 24h](#)


**45%**







# UChicago Network Statistics

- Using CollectD's SNMP plugin to gather data from our switches
- Switches are tricky.
  - Interface counters measure octets and can roll-over to 0
  - nonNegativeDerivative and scaling functions are needed

# UChicago Network Statistics




# UChicago Network Statistics



# HTCondor Statistics

- We monitor HTCondor for both MWT2 and for our OSG Connect service.
- For MWT2, interested in job efficiencies, etc.
- For OSG Connect, interested in behavior of the submit box, etc.


# HTCondor Statistics



# HTCondor Statistics



# HTCondor Statistics



# Future directions

- Graphite has a mechanism for tagging events
  - Basically a vertical bar on the graph, viewable in Grafana
  - e.g., “SRM errors increased due to switch problems”
- Grafana’s templated dashboards to create per-user or per-site metrics



# Finally..

- Feel free to explore our Grafana instance:
  - <http://grafana.mwt2.org:3000/>
- Code: <https://github.com/DHTC-Tools/grafana-collectors>
  - (n.b., this is highly disorganized at the moment. )

# Thanks

- Questions?